Spider and the Flies : Focused Crawling on Tumblr to Detect Hate Promoting Communities

Swati Agarwal
Indraprastha Institute of Information Technology, Delhi (IIIT-D), India
swatia@iiitd.ac.in

Ashish Sureka
ABB Corporate Research Labs
Bangalore, India
ashish.sureka@in.abb.com

arXiv:1603.09164v1 [cs.SI] 30 Mar 2016

ABSTRACT

Tumblr is one of the largest and most popular microblogging website on the Internet. Studies shows that due to high reachability among viewers, low publication barriers and social networking connectivity, microblogging websites are being misused as a platform to post hateful speech and recruiting new members by existing extremist groups. Manual identification of such posts and communities is overwhelmingly impractical due to large amount of posts and blogs being published every day. We propose a topic based web crawler primarily consisting of multiple phases: training a text classifier model consisting examples of only hate promoting users, extracting posts of an unknown tumblr microblogger, classifying hate promoting bloggers based on their activity feeds, crawling through the external links to other bloggers and performing a social network analysis on connected extremist bloggers. To investigate the effectiveness of our approach, we conduct experiments on large real world dataset. Experimental results reveals that the proposed approach is an effective method and has an F-score of 0.80. We apply social network analysis based techniques and identify influential and core bloggers in a community.

Keywords

Hate and Extremism Detection, Information Retrieval, Microblogging, Mining User Generated Content, Online Radicalization, Social Network Analysis, Topical Crawler

1. INTRODUCTION

Tumblr is the second largest microblogging platform, has gained phenomenal momentum recently. It is widely used by fandoms: communities of users having similar interests in various TV shows and movies [15]. Therefore, it is especially popular among young generation users and provides them a platform to discuss daily events. They communicate by blogging and publishing GIF images as their reactions and emotions on several topics [5]. According to Tumblr statistics 2015 [1], over 219 million blogs are registered on Tumblr and 420 million are the active users. 80 million posts are being published everyday, while the number of new blogs and subscriptions are 0.1 million and 45 thousands respectively.

Tumblr is also posed as a social networking website that facilitates users to easily connect to each other by following other users and blogs without having a mutual confirmation. Bloggers can also communicate via direct messages that can be sent privately or can be posted publicly using 'ask box'. It facilitates bloggers to send these messages anonymously if they don't want to reveal their Tumblr identity [13]. Similar to other social networking websites, Tumblr has very low publication barriers. A blogger can publish a new post and can re-blog an existing public post which is automatically broadcasted to it's followers unless it is enabled as a private post [7]. The type of posts can be chosen among seven different categories including multi-media and other content: text, quote, link, photo, audio, video and URL. Unlike Twitter, Tumblr has no limit on the length of textual posts. Similar to hashtags in Twitter, there are separated tags associated with the blog content that make a post easier to be searchable on Tumblr [2]. Tumblr also allows users to update their other connected social networking profiles when something is posted.

The simplicity of navigation, high reachability across wide range of viewers, low publication barriers, social networking and anonymity has led users to misuse Tumblr in several ways. Previous studies shows that these features of Tumblr are exclusive factors to gain the attention of modern extremist groups [4][16]. This is because Tumblr provides every kind of multimedia posts which is a great medium to share your views with your target audiences. These groups form their own communities that share a common propaganda. They post rude comments against a religion to express their hatred and spread extremist content. Social networking facilitates these groups to recruit more people to promote their beliefs and ideology among global audiences [8][12]. Figure 1 illustrates a concrete example of various types of hate promoting posts and their associated tags on Tumblr. The number of notes shows the number of times that post has been liked and re-blogged by other blogs.

Online-radicalization and posting hateful speech is a crime against the humanity and mainstream morality; it has a major impact on society [8]. Presence of such extremist content on social media is a concern for law enforcement and intelligence agencies to stop such promotion in country as it poses the threat to the society [3][9]. It also degrades the reputation of the website and therefore is a concern for website moderators to identify and remove such communities. Due to the dynamic nature of website, automatic identification of extremist posts and bloggers is a technically challenging problem [17]. Tumblr is a large repository of text, pictures and other multimedia content which makes it impractical to search for every hate promoting post using keyword based flagging. The textual posts are user generated data that con-

[1]https://www.tumblr.com/about

[2]https://www.tumblr.com/docs/en/using_tags

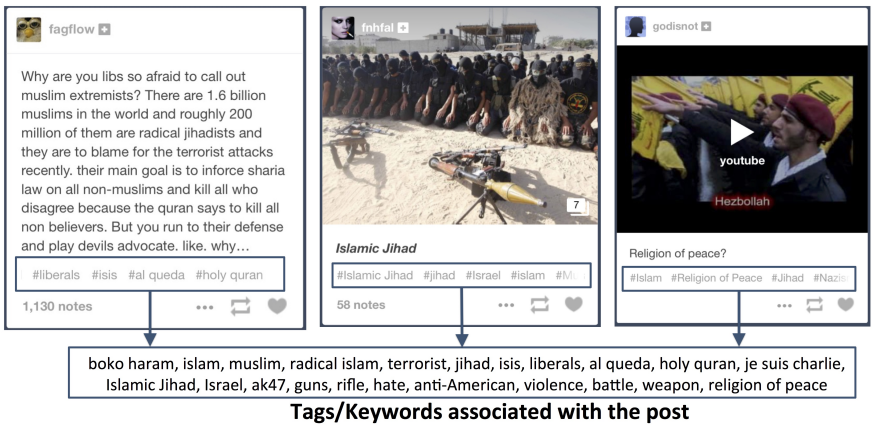
fagflow
Why are you libs so afraid to call out muslim extremists? There are 1.6 billion muslims in the world and roughly 200 million of them are radical jihadists and they are to blame for the terrorist attacks recently. their main goal is to inforce sharia law on all non-muslims and kill all who disagree because the quran says to kill all non believers. But you run to their defense and play devils advocate. like. why…
#liberals #isis #al queda #holy quran
1,130 notes
Inhfal
Islamic Jihad
#Islamic Jihad #jihad #Israel #islam
58 notes
godisnot
youtube
Hezbollah
Religion of peace?
#Islam #Religion of Peace #Jihad
boko haram, islam, muslim, radical islam, terrorist, jihad, isis, liberals, al queda, holy quran, je suis charlie, Islamic Jihad, Israel, ak47, guns, rifle, hate, anti-American, violence, battle, weapon, religion of peace
Tags/Keywords associated with the post

Figure 1: A Concrete Example of Different Types (Text, Video, Audio) of Hate Promoting Post on Tumblr

tain noisy content such as spelling, grammatical mistakes, presence of internet slangs and abbreviations. Presence of low quality content in contextual metadata poses technical challenges to text mining and linguistic analysis [3][13]. The work presented in this paper is motivated by the need of investigating solutions to counter and combat the online extremism on Tumblr.

The research aim of the work presented in this paper is the following:

1. To investigate the application of topical crawling based algorithm for retrieving hate promoting bloggers on Tumblr. Our aim is to examine the effectiveness of random walk in social network graph graph traversal and measuring its performance.
2. To investigate the effectiveness of contextual metadata such as content of the body, tags and caption or title of a post for computing the similarity between nodes in graph traversal. To examine the effectiveness of *reblogging* and *like* on a post as the links between two bloggers.
3. To conduct experiments on large real world dataset and demonstrate the effectiveness of proposed approach in order to locate virtual and hidden communities of hate and extremism promoting bloggers and apply Social Network Analysis based techniques to locate central and influential users.

2. LITERATURE SURVEY

In this section, we discuss closely related work to the study presented in this paper. Based on our review of existing work, we observe that most of the researches for detecting online radicalization are performed on Twitter, YouTube and various discussion forums [1]. We conduct a literature survey in the area of identifying hate promoting communities on social networking websites and short text classification of Tumblr microblog . O'Callaghan et. al. [14] describe an approach to identify extreme right communities on multiple social networking websites. They use Twitter as a possible gateway to locate these communities in a wider network and track dynamic communities. They perform a case study using two different datasets to investigate English and German language communities. They implement a heterogeneous network within a homogeneous network and use four different social networking platforms (Twitter accounts, Facebook profiles, YouTube channels and all other websites) as extreme right entities or peers and edges are the possible interactions among these accounts.

Mahmood S. [12] describes several mechanisms that can be useful in order to detect presence of terrorists on social networking websites by analyzing their activity feeds. They use Google search and monitor terror attack using keyword-based flagging mechanism. They monitor sentiments and opinions of users following several terrorism groups on online social networks and propose a counter-terrorism mechanism to identify those users who are more likely to commit a violent act of terror. They also discuss honeypots and counter-propaganda techniques that can be used to rehabilitate radicalized users back to normal users. The disadvantage of keyword based flagging approach is the large number of false alarms. David and Morcelli [8] present a keyword based search to detect several criminal organizations and gangs on Twitter & Facebook. They discuss a study of analyzing the presence of organized crime and how these gangs use social media platforms to recruit new members, broadcast their messages and coordinate their illegal activities on web 2.0. They perform a qualitative analysis on 28 groups and compare their organized crime between 2010 and 2011 on Facebook.

Agarwal et. al. propose a one-class classification model to identify hate and extremism promoting tweets [16]. They conducted a case study on Jihad and identified several linguistic and stylistic features from free form text such as presence of war, religious, negative emotions and offensive terms. They conduct experiments on large real world dataset and demonstrate a correlation between hate promoting tweets and discriminatory features. They also perform a leave-p-out strategy to examine the influence of each feature on classification model.

In context to existing work, the study presented in this paper makes the following unique contributions extending our previous work [2]:

1. We present an application of topical crawler based approach for locating extremism promoting bloggers on Tumblr. While there has been work done in the area of topical based crawling of social media platforms, to the best of our knowledge this paper is the first study on topical crawling for navigating connections between Tumblr bloggers.

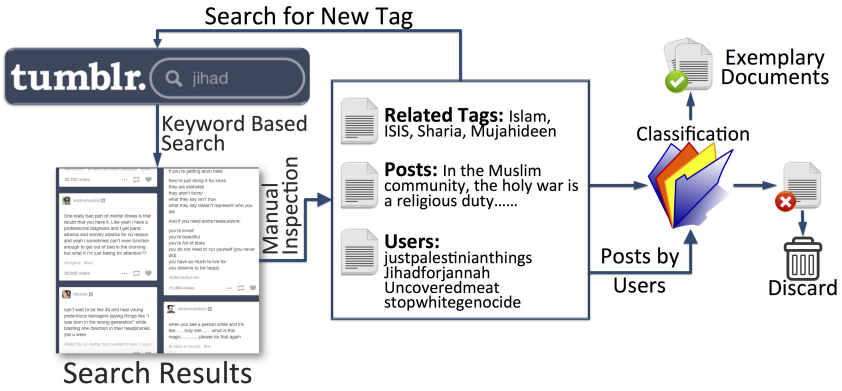

Figure 2: Flow Sequence of Exemplary Data Collection Process

2. We conduct experiments on large real world dataset to demonstrate the effectiveness of one class classifier and filtering hate promoting blog posts (text). We retrieve Tumblr blogger profiles and their links with other hate promoting bloggers and apply Social Network Analysis to locate strongly connected communities and core bloggers.

3. EXPERIMENTAL SETUP

We conduct our experiments on an open source and real time data extracted from Tumblr micro-blogging website. In a social networking website, a topical crawler extracts the external link to a profile and returns the nodes that are relevant to a defined topic. We define the relevance of a node based on extent of similarity of it's activity feeds and training document. Topical crawler learns the features and characteristics from these training documents and classify a profile to be relevant. Figure 2 illustrates the general framework to obtain these documents. As shown in Figure 2, we implement a bootstrapping methodology to collect the training samples. We perform a manual search on Tumblr and create a lexicon of popular and commonly used tags associated with hate promoting posts. Figure 3 shows a word cloud of such terms. To collect our training samples, we perform a keyword (search tag) based flagging and extract their associated textual posts. We also acquire the related tags and the linked profiles (users who made these posts). We expand our list of keywords by extracting associated tags from these posts and their related tags. We run this framework iteratively until we get a reasonable number of exemplary documents (400 training samples). As mentioned above, we train our classifier for only hate promoting users. Therefore, the training documents contain the content and caption of only extremist posts.

We use these linked bloggers and posts to compute the threshold value for language modeling. We take a sample of 30 bloggers and compare their posts with the exemplary documents. For each blogger we get a relevance score. To compute the threshold value for similarity computation we take an average of these scores. Figure 4 illustrates the relevance score statistics of each blogger (Sorted in increasing order). We notice that 80% of the bloggers have relevance score between -2.7 and -1.5. We take average (turns out to be -2.58) of these scores to avoid the under-fitting and over-fitting of bloggers during classification.

Figure 3: A Word Cloud of Key Terms Commonly Used by Extremist Bloggers

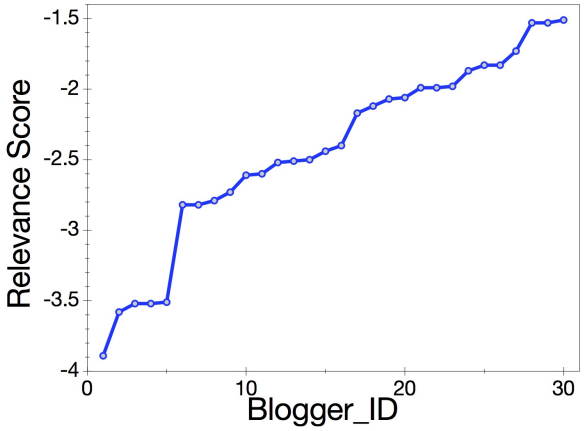

Figure 4: Illustrating the Relevance Score Statistics of Positive Class Bloggers

4. RESEARCH METHODOLOGY

In this section, we present the general research framework and methodology of proposed approach for classifying extremist bloggers on Tumblr (refer to Figure 5). The proposed approach is an iterative multi-step that uses a hate promoting blogger as a seed channel and results a connected graph where nodes represents the extremist bloggers and links represents the relation between two bloggers (like and re-blog). As shown in Figure 5, proposed framework is a multi-step process primarily consists of four phases: i) ex-

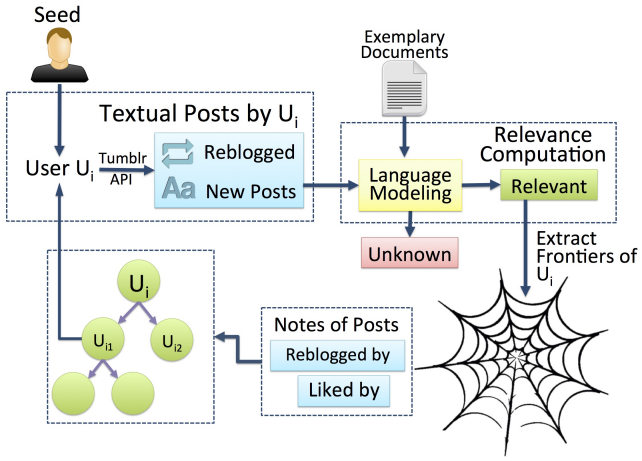

Figure 5: Proposed Architecture for Extremist Community Detection. Source: Agarwal et. al. [2]

traction of activity feeds of a blogger, ii) training a text classification model and filtering hate promoting and unknown bloggers, iii) navigating through external links to bloggers and extracting linked frontiers and iv) traversing through spider network for selecting next blogger. In Phase 1, we use a positive blogger U_i (annotated as hate promoting during manual inspection) called as 'seed'. We extract n number of textual posts (either re-blogged or newly posted by user)of U_i by using Tumblr API [3]. We further use Jsoup Java library [4] to extract the content and caption of these posts. Tumblr allows users to post content in multiple languages. However, our focus of this paper is to mine only English language posts. Therefore, we perform data-preprocessing on all extracted posts and by using Java language detection library [5], we filter all non-English language posts. In Phase 2, we train our classification model over training samples (refer to Section 3). We perform character level n-gram language modeling[6] on English language posts and compute their extent of similarity against training samples. We classify a blogger as hate promoting based on the relevance score and computed threshold value (refer to Section 3).

If a channel is classified as hate promoting or relevant, we further proceed to Phase 3 and extract the notes information for each posts (collected in phase 1). Notes in a Tumblr post contains the information about bloggers who liked or re-blogged a post. These user hits on a post indicates the similar interest among bloggers who may or may not be direct followers of each other. We extract the Tumblr ids of profiles from notes for the following reasons: i) due to privacy policy, Tumblr does not allow users/developers to extract the followers list unless the list is public and ii) Tumblr facilitates users to track any number of search tags or keywords. Whenever a new post is published on Tumblr containing any of these tags, it appears on the dashboard of user and a blogger no longer need to follow the original poster. We manage a queue of all extracted bloggers and traverse through the network using Random Walk algorithm. We use uniform distribution to select next blogger and extract it's frontiers. We extract these bloggers until the graph converges without re-visiting a blogger. The proposed framework results into a connected graph that represents a Tumblr network. We perform social network analysis on the output graph to locate hidden virtual groups and extremist bloggers playing major roles in community.

[3]https://www.tumblr.com/docs/en/api/v2
[4]http://jsoup.org/apidocs/
[5]https://code.google.com/p/language-detection/
[6]http://alias-i.com/lingpipe/index.html

Algorithm 1: Extracting Textual Posts on Tumblr

Data: User U, Consumer Key C_k, Consumer Secret C_s, Search Tag *tag_name*
Result: Text based posts made by User U or associated with tag *tag_name*

```
1  SetParameters()
2  TaggedPost()
3  BloggerPost()
   Generate URL of post to fetch post content and caption
4  for all postP ∈ Posts do
5      Slug=P.getSlug()
6      id=P.getID()
7      URL="http://blog_name.tumblr.com/post/id/slug"
8      Document=Jsoup.connect(URL).get()
9      post_content=Document.getDescription()
10     post_caption=Document.getTitle()
   end
   SetParameters() {
11 Authenticate the client via API Keys C_k and C_s
12 params.put("type", "text")
13 params.put("filter", "text")
14 params.put("reblog info", true)
15 params.put("notes info", true) }
   TaggedPost() {
16 Posts = client.tagged(tag_name, params) }
   BloggerPost(){
17 Posts = client.tagged(tag_name, params) }
```

5. SOLUTION IMPLEMENTATION

A topical crawler starts from a seed node, traverses in a graph navigating through some links and returns all relevant nodes to a given topic. In proposed solution approach we divide our problem into three sub-problems. First we classify the given seed node S as hate promoting or unknown according to the published post (originally posted by blogger or re-blogged from other Tumblr users). Second, if the node is relevant then we extend this node into it's frontiers and it further leads us to more hate promoting bloggers. In third sub-problem, we perform topical crawling on Tumblr network and use random walk algorithm to traverse along the graph.

5.1 Retrieval of Published Posts

Algorithm 1 describes the method to search Tumblr posts using keyword based flagging and extraction of posts published by a given blogger. The work presented in this paper focus on mining textual metadata on Tumblr therefore we set a few parameters and extract only text based posts for further analysis. For each blogger we set the limit of 100 posts published recently. Function `SetParameters()` (steps 12 and 13) filters the search results and displays only the tex-

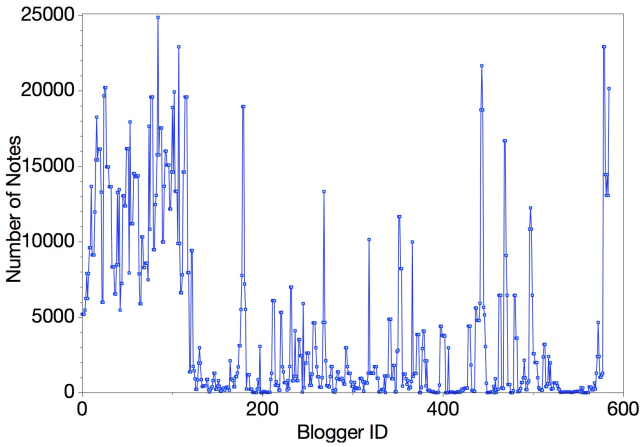

Figure 6: Illustrating the Number of Notes For Each Blogger Traversed in Topical Crawler

tual posts (quote, chat, text, url). Function `TaggedPost()` with given parameters search for text posts that exclusively contain given tag name. `BloggerPost()` fetches the textual posts published by given blogger ID. Both the functions make a Tumblr API request to fetch these data. Function `ExtractPost()` filters the response and extract body content & caption of each post. Tumblr API allows us to only extract the summary of large posts. Therefore we use HTML parsing for extracting the whole message in blog post. In steps 4 to 7, we generate the URL from post summary and id to fetch the remaining post details. ID is a unique identifier of Tumblr posts and slug is a short text summary of that post which is appended in the end of every URL. We invoke this URL using Jsoup library and parse the HTML document to get the post content.

Algorithm 2: Extracting Frontiers of a Given Blog

```
Data: Blogger U
Result: Frontiers F < name, type > of U
ExtractFrontiersU{
1  SetParameters()
2  Posts=BloggerPost(U)
3  for all postP ∈ Posts do
4      Notes=P.getNotes()
5      for all NoteN ∈ Notes do
6          Linked_Blog_Name=N.getBlogName()
7          Note_Type=N.getType()        Liked or Reblog
8          Frontiers F.add(Linked_Blog_Name,
           Note_Type)
       end
   end
   }
```

5.2 Retrieval of External Links to Bloggers

Algorithm 2 describes the steps to extract frontiers of a given node U. Due to the privacy policies, Tumblr API does not allow developers to extract subscriptions and followers of a Tumblr user. The link between two bloggers indicates the similar interest between them so that number of frontiers vary for every post published by a blogger. For each user,

Algorithm 3: Graph Traversal Using Random Walk Algorithm

```
   Data: S, th, N_g, S_g, W_g, D_e
   Result: Directed Graph G
1  SetParameters()
2  U_i = S, F.add(S)
   TopicalCrawlerS{
3  while (graphsize < S_g OR F.size > 0) do
4      Posts=ExtractPost(U_i)
5      Relevance_Score = LanguageModeling(D_e, Posts,
       N_g)
6      if (score >th) then
7          Linked_Users=ExtractFrontiers(U_i)
8          ProcessedNodes PN.add(U_i)
9          for all LU ∈ Linked_Users do
10             if (!(F.contains(LU) AND
               (PN.contains(LU))) then
11                 F.add(LU)
               else
12                 Discard the node LU
               end
           end
       else
13         Discard the node U_i
       end
14     Compute the Markov Chain over graph G
15     New_Blogger= node with maximum probability in
       Markov chain array
16     F.remove(New_Blogger)
17     TopicalCrawler(New_Blogger)
   end
   }
```

we extract 25 bloggers for each relation i.e. users who have liked and re-blogged that post recently. If a blogger B_1 reblog and as well as likes a post published by another blogger B_2 then in the graph G, we create an edge with both labels i.e. (B_1, B_2, <like, re-blog>). To avoid the redundancy we extract one more frontier who have either liked or reblogged the post recently. To extract the linked bloggers of a Tumblr user we first need to extract the posts made by U. We can extract notes information only when notes and re-blog information parameters are set to be true (refer to steps 14 and 15 of Algorithm 1). As described in Algorithm 2, in step 4, we extract notes for each textual post (hate promoting) made by User U. In steps 5 to 8, we extract the name of unique bloggers who liked and re-blogged the post P. In step 8, F represents the list of frontiers and relation of U with each frontier. We maintain a list of all processed bloggers and the number of hit counts on their recent 100 posts. These number of notes varied from 0 to 25K therefore we perform smoothing on data points and plot median of these values. Figure 6 shows the statistics of number of notes collected on 100 posts of each blogger extracted during topical crawler. Figure 6 reveals that overall number of hit counts (number of reblog and like) for extremism promoting users is very high. These hit counts reveals the popularity of extremist content and the number of viewers connected to such bloggers.

5.3 Topical Crawler Using Random Walk

Algorithm 3 describes the proposed crawler for locating a group of hidden extremist bloggers on Tumblr. The goal of this algorithm is to compare each blogger against training examples and then connecting all positive class (hate promoting or relevant) bloggers. Algorithm 3 takes several inputs: seed blogger (positive class user) S, size of the graph S_g i.e. maximum number of nodes in a graph, width of the graph W_g i.e. the maximum number of frontiers or adjacency nodes for each blogger, a set of exemplary documents D_e, threshold th and n-gram value N_g for relevance computation. We create a list of 30 positive class bloggers extracted during experimental setup (refer to section 3) and compute their relevance score against the exemplary documents. We take an average of these scores and compute the threshold value for language modeling. We use n-gram language modeling (N_g=3) to build our statistical model. Algorithm 3 is a recursive process that results into a cyclic directed graph. We run this algorithm until we get a graph of size S (1000 bloggers) or there is no node left in the queue for further extension. We perform a self-avoiding random walk that means we make sure a node is never being re-visited. If a node re-appears in the frontiers list then there are two possibilities: 1) the frontier has already been processed (extended or discarded based upon the relevance score- Steps 4 to 7). If it exists in the processed nodes list then we create a directed edge between the node and it's parent and avoid further extension. 2) If the re-appearing node is in frontiers list and is not yet processed, we created a directed edge in the graph and continue the traversal.

The topical crawler is a recursive process that adds and removes nodes after each iteration. The resultant graph is dynamic and not irreducible that means given a graph $G(V, E)$, if there is a directed edge between two nodes u and v, it is not necessary that there exists a directed path from v to u. Consider that object (topical crawler) processed node i at time $t-1$. In the next iteration object moves to an adjacency node of i. The probability that object moves to node j at time t is $\frac{1}{d_i^+}$ when there exists a direct edge from i to j. M_{ij}=$\frac{1}{d_i^+}$ denotes the probability to reach from i to j in one step where d_i^+ is the out-degree of node i. Therefore we can define:

$$M_{i,j} = \left\{ \begin{array}{ll} \frac{1}{d_i^+}, & if\ (i,j)\ is\ an\ edge\ in\ digraph\ G \\ 0, & otherwise \end{array} \right\} \quad (1)$$

Therefore for each vertex i, the sum of the probability to traverse an adjacency node of i is 1.

$$\forall i \sum_{j \in A(i)} M_{ij} = 1 \quad (2)$$

Where, A(i) denotes the list of adjacency nodes i. In random walk on graph G topical crawler traverse along the nodes according to the probability of M_{ij}. Graph G is a dynamic social networking graph, therefore we compute a Markov chain M after each iteration and compute the probability matrix over graph G. Markov chain is a random process where the probability distribution of node j depends on the current state of matrix. The probability matrix M^k gives us a picture of graph G after k iterations of topical crawler. Using this matrix, we compute the probability distribution P that object moves to a particular vertex. P^k is the probability distribution of a node j after k iteration then probability of i to be traversed in $k+1^{th}$ iteration is the following:

$$P^{k+1} = P^k.Mwhere, P^k = P^0 * M^k \quad (3)$$

Where P^0 is the initial distribution fixed for the seed node.

Table 1: Confusion Matrix and Accuracy Results for Unary Classification Performed During Topical Crawler. Source: Agarwal et. al. [2]

(a) Confusion Matrix

		Predicted	
		Positive	**Unknown**
Actual	**Positive**	290	45
	Unknown	92	173

(b) Accuracy Results

Precision	**Recall**	**F-Score**	**Accuracy**
0.75	0.86	0.80	0.77

6. EXPERIMENTAL RESULTS

6.1 Topical Crawler Results

We hired 30 graduate students from different departments as volunteers to annotate all 600 bloggers processed or traversed during the topical crawling. We provided them guidelines for annotation and to remove the bias, we performed horizontal and vertical partitioning on the bloggers' dataset. We divided annotators into 10 different groups, 3 members each. We asked each group to to annotate 60 bloggers. Therefore, we got 3 reviews for each blogger. We used majority voting approach (a blogger labeled as X by at least two annotators) for annotating each blogger as hate promoting or unknown. We compute the effectiveness of our classification by using precision, recall and f-measure as accuracy matrices. Table 1 shows the confusion matrix and accuracy results for unary classification performed during graph traversal. Table 1(a) reveals that among 600 bloggers, our model classifies 382(290+92) bloggers as hate promoting and 218(173+45) bloggers as unknown with a misclassification of 13% and 34% in predicting extremist and unknown bloggers respectively. Table 1(b) shows the standard information retrieval matrices for accuracy results. our results shows that the both the precision and recall are high, as it is important to reduce the number of false alarms and also not to miss an extremist blogger in order to locate their communities. We use F-score as the accuracy metrics for our classifier and the results reveal that we are able to predict extremist blogger with an f-score of 80%.

6.2 Social Network Analysis

We perform social network analysis on topical crawler's network resulted into a directed graph $G(V, E)$, where V represent a set of Tumblr bloggers accounts and E represent a directed edge between two bloggers. We define this edge as a relation having two labels 'posts liked by' and 'post reblogged by'. To examine the effectiveness of these relations we generate two independent networks exclusively for 'like' and 're-blog' links between bloggers. Figure 7 illustrates the

Table 2: Illustrating The Network Level Measurements for Topical Crawler.Dia= Network Diameter, Mod= Modularity, ACC= Average Clustering Coefficient, IBC= In- Betweenness Centrality, CC= In- Closeness Centrality, #SCC= Number of Strongly Connected Components.

Graph	#Nodes	#Edges	Dia	#SCC	#ACC	#Mod	IBC	ICC
TC	382	275	4	137	0.026	12.00	11.36	0.20
LB	27	60	1	21	0.0231	1.307	0	0.38
RB	355	215	6	185	0.021	7.01	6.284	0.40

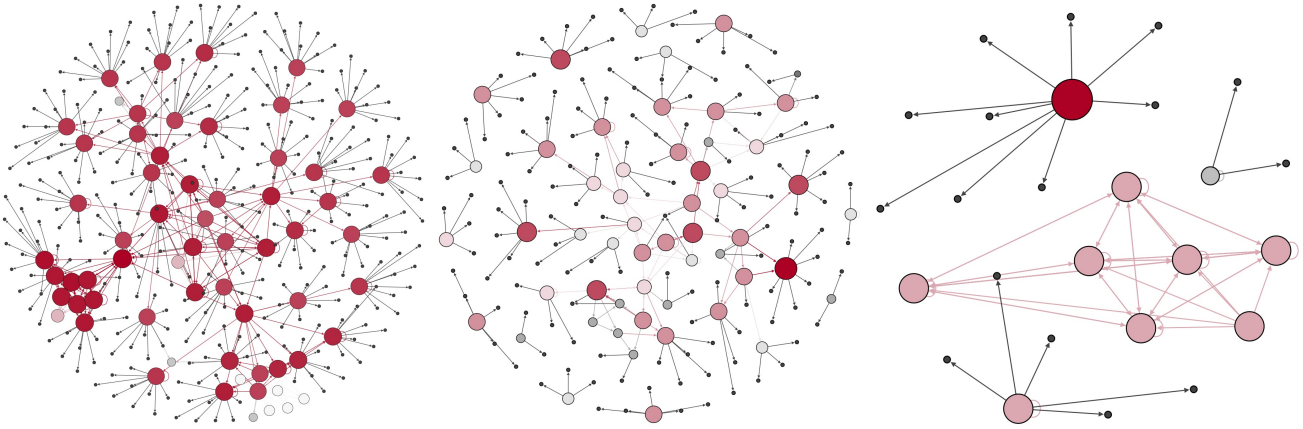

(a) (b) (c)

Figure 7: Cluster Representation of Social Network Graphs- Topical Crawler using Random Walk (a), 'Posts Liked by' (b) and 'Posts Re-blogged By' (c)

representations of these networks. In each graph, size of the node is directionally proportional to its out-degree. A node with maximum number of adjacency vertices is biggest in size. Colours in the graph represents the clusters of nodes having similar properties. Here, we define the similarity measure as the ratio of out-degree and in-degree.

We also perform several network level measurements on these graphs. Table 2 reveals that re-blogging is a strong indicator of links between two Tumblr profiles. Here, we observe that the graphs generated for topical crawler and re-blogging link have same pattern in network measurements. Both graphs are dense (also evident from the Figure 7) and have higher modularity in comparison to the graph created for 'liked' link. Table 2 and Figure 7 also reveal that by navigating through re-blogging links we can locate large number of connected components in a extreme right community. While following 'like' as a link we are able to detect small number of connected blogs. Though as illustrated in Figure 7(b), we can not completely avoid this feature as a set of blogs extracted using this link are irreducible. Table 2 also shows that the graph created for 'like' relation has slightly larger value for average clustering co-efficient. This is because the number of nodes in the graph is very less and a major set of these nodes is strongly connected. Higher value of In-between centrality shows the presence of bloggers who are being watched by a large number of users. As the Figure 7(c) shows there are many users which are not directly connected to each other (shown in red colour) but has a huge network of common bloggers. These disjoint bloggers are two or three hop away and are connected via other bloggers (having second largest number of adjacency nodes). These bloggers are connected with maximum number of other bloggers present in the graph and has a wide spread network in extreme right communities. These nodes have the maximum closeness centrality and play central role in the community. Nodes represented as black dots have minimum number of out-degree. They don't have a directed path to the central users or original source of extremist posts. Based upon our study we find these bloggers to be the target audiences who share these posts in their own network. These users are very crucial for such communities though they don't play a major role in the network.

7. CONCLUSIONS

In this paper, we perform a case study on Jihadist groups and locate their existing extreme right communities on Tumblr. We conduct experiments on real world dataset and use topical crawler based approach to collect textual data (published posts) from Tumblr users. We perform one class classification and identify hate promoting bloggers according to the content present in their posts. We use random walk algorithm for graph traversal and extract exclusive links to these bloggers. We conclude that by performing social networking analysis on a graph (vertices are the Tumblr bloggers and edges are the links among these bloggers: re-blog and like) we are able to uncover hidden virtual communities of extremist bloggers with an accuracy of 77%. We compute various centrality measures to locate the influential bloggers playing major roles in extremist groups. We also investigate the effectiveness of link features (likes and re-blogs) in order to find the communities. Our results reveal re-blogging is a strong indicator and a discriminatory feature to mine strongly connected communities on Tumblr.

We perform a manual inspection on Tumblr and perform

a characterization on several hate promoting posts. Our study reveals that these posts are very much popular among extremist bloggers and get large number of hits. These posts are published targeting some specific audiences. Keywords present in the blog content, tags associated with post and comments by other bloggers are clear evidence of hate promotion among their viewers.

8. REFERENCES

[1] S. Agarwal and A. Sureka. Applying social media intelligence for predicting and identifying on-line radicalization and civil unrest oriented threats. *arXiv preprint arXiv:1511.06858*, 2015.

[2] S. Agarwal and A. Sureka. A topical crawler for uncovering hidden communities of extremist micro-bloggers on tumblr. In *5th Workshop on Making Sense of Microposts (MICROPOSTS)*, 2015.

[3] S. Agarwal, A. Sureka, and V. Goyal. *Big Data Analytics: 4th International Conference, BDA 2015, Hyderabad, India, December 15-18, 2015, Proceedings*, chapter Open Source Social Media Analytics for Intelligence and Security Informatics Applications, pages 21–37. Springer International Publishing, 2015.

[4] R. A. Bates and M. Mooney. Psychological operations and terrorism: The digital domain. *The Journal of Public and Professional Sociology*, 6(1):2, 2014.

[5] E. Bourlai and S. C. Herring. Multimodal communication on tumblr: "i have so many feels!". In *Proceedings of the 2014 ACM Conference on Web Science*, WebSci '14, pages 171–175, New York, NY, USA, 2014. ACM.

[6] R. Bury, R. Deller, A. Greenwood, and B. Jones. From usenet to tumblr: The changing role of social media. *Participations*, 10(1):299–318, 2013.

[7] Y. Chang, L. Tang, Y. Inagaki, and Y. Liu. What is tumblr: A statistical overview and comparison. *CoRR*, abs/1403.5206, 2014.

[8] D. Décary-Hétu and C. Morselli. Gang presence in social network sites. *International Journal of Cyber Criminology*, 5(2), 2011.

[9] M. R. DeVore. Exploring the iran-hezbollah relationship: A case study of how state sponsorship affects terrorist group decision-making. *Perspectives on Terrorism*, 6(4-5), 2012.

[10] J. Gies and S. Martino. Uncovering ed: A qualitative analysis of personal blogs managed by individuals with eating disorders. *The Qualitative Report*, 19(29):1–15, 2014.

[11] D. Jurgens. That's what friends are for: Inferring location in online social media platforms based on social relationships. In *ICWSM*, 2013.

[12] S. Mahmood. Online social networks: The overt and covert communication channels for terrorists and beyond. In *Homeland Security (HST), 2012 IEEE Conference on Technologies for*, pages 574–579. IEEE, 2012.

[13] E. Marquart. Microblog sensation: The growing popularity of tumblr. *3PM Journal of Digital Research and Publishing, Session*, 2:70–75, 2010.

[14] D. O'Callaghan, D. Greene, M. Conway, J. Carthy, and P. Cunningham. Uncovering the wider structure of extreme right communities spanning popular online networks. In *Proceedings of the 5th Annual ACM Web Science Conference*, pages 276–285. ACM, 2013.

[15] L. Renwick. Audience research project: Tumblr study group research 'how do 'fandoms' on tumblr react to new media content?'. *Enquiry-The ACES Journal of Undergraduate Research*, 4, 2014.

[16] A. Sureka and S. Agarwal. Learning to classify hate and extremism promoting tweets. In *Intelligence and Security Informatics Conference (JISIC), 2014 IEEE Joint*, pages 320–320. IEEE, 2014.

[17] A. S. Wilner and C.-J. Dubouloz. Transformative radicalization: Applying learning theory to islamist radicalization. *Studies in Conflict & Terrorism*, 34(5):418–438, 2011.

[18] J. Xu, R. Compton, T.-C. Lu, and D. Allen. Rolling through tumblr: Characterizing behavioral patterns of the microblogging platform. In *Proceedings of the 2014 ACM Conference on Web Science*, WebSci '14, pages 13–22, New York, NY, USA, 2014. ACM.

[19] J. Xu, T.-C. Lu, R. Compton, and D. Allen. Civil unrest prediction: A tumblr-based exploration. In *Social Computing, Behavioral-Cultural Modeling and Prediction*, pages 403–411. Springer, 2014.